ESDS Technical Report

The Measurement of Poverty, Welfare, and Distribution: Basic Issues

July 1997

Michael Shea
Economic & Social Data Service (PPC/CDIE/DI)

This report was prepared for the Office of Emerging Markets in the Global Bureau's Center for Economic Growth and Agricultural Development, under contract # AEP-1016-006

ESDS Technical Report

The Measurement of Poverty, Welfare, and Distribution: Basic Issues

Michael Shea
ESDS
July 1997

There are three distinct -- yet related -- concepts used to assess the social and economic condition of a country or region: poverty, welfare, and distribution. Briefly defined, they are:

1. Poverty - the lack of command over basic consumption needs
2. Welfare - the level of a standard-of-living or well-being measure
3. Distribution - a comparison of income or wealth among different groups

Although a correlation between these concepts can often be established empirically, it is misleading to use one as if it were identical to another. For example, data that shows income has become more unequally distributed ***is not*** evidence that poverty has increased. Similarly, data showing growth in the median household income ***is not*** proof that poverty has declined.

The following sections describe commonly used measures of the three concepts.

Poverty

The first step in measurement of poverty is to determine the level at which consumption of or access to goods and services is considered to be "basic" or "minimally acceptable."[1] At that point, a line is drawn and individuals or households below the line are classified as poor.[2] Ravallion and Sen (1996[3]) identify two broad categories of methods:

< the cost-of-basic-needs approach (CBN) and
< the food-energy-intake approach (FEI).

The CBN method assembles a basket of goods and services that are deemed to meet a minimal standard of living, appropriate to the society in which the households live. Critics note that the choice of goods in the "basic" basket is seen as arbitrary and that the price data needed to calculate the poverty line are often incomplete. Alternatively, the FEI approach determines the income or consumption level at which a household is typically just sufficient for energy needs. This approach eliminates the need for any price data, as well as any arbitrariness with respect to non-food consumption.[4] The line can be drawn by identifying households (from household survey data) which intake "just enough" food, and then estimating their mean income or consumption.[5] A slightly different way to draw the line would be to run regressions between food energy intake and consumption expenditures, and then use the coefficients to estimate the poverty line.

Comparing Poverty Lines in Different Countries

Poverty lines that are drawn will vary from country to country. In some countries, basic telephone service is seen as a necessity, but not in others. With some uncertainty over the "accuracy" of any poverty line drawn, a series of poverty lines can be adopted. In the *Social Indicators of Development*, for example, the World Bank reports a "lower" and an "upper" poverty line. Individuals or households falling below the "lower" poverty line can be said to be in "extreme poverty." In that World Bank database, the poverty lines are still country-specific, and vary in value from country to country. (Note also that more than ninety percent of the countries in the database do not have poverty lines available at all. Relative scarcity is a cardinal feature of poverty data, at least for now.)

As many analysts have noted, the "country-specific" poverty line is typically set at an absolute level within that country, *i.e.*, it has a constant value at different points in time and is not adjusted when average incomes rise over the short-term. However, the level of that poverty line -- when converted in a common unit, *e.g.*, at current exchange rates into U.S. dollars or by purchasing-power-parity adjusted dollars -- has a marked tendency to be higher if the country is in a higher income category. Thus, data published in 1995 indicated that the upper poverty line was sixteen dollars in Lao PDR, three hundred dollars in the Philippines, and nearly thirteen hundred dollars in Georgia. (All numbers reported in local currency units and then converted into U.S. dollars at prevailing exchange rates.)

In order to make cross-country comparisons of poverty data *that classify individuals as*

poor by the same yardstick, the World Bank developed its absolute poverty line of $370. (The dollar units are adjusted for "purchasing power parities," a method which estimates equivalent values of local currencies in different countries.[6]) This poverty line is sometimes referred to as the "dollar-a-day" poverty line. Using the dollar-a-day approach, the Bank estimated that there were 1.1 billion poor people in the developing world in 1990, or about thirty percent of the total population.[7] In addition, the Bank refers to a lower poverty line, or a measure of extreme poverty, put at $275 per person, per year.[8]

Although this global benchmark is useful from a trends analysis or regional aggregate perspective, it is less appropriate as a policy evaluation measure in the country-specific context. For example, the "dollar-a-day" poverty line approach estimated that two-thirds of the population in India was in poverty in 1990, but only about five percent in Thailand the same year.[9] Does this mean that the Thai government can dismiss poverty as a negligible problem? Only if their policy goal is to ensure that every citizen is at least a little better off than the poor in India. In fact, a higher poverty line should be, and is, used in Thailand, where the per capita income of the total population is several times as high as in India.

Key Features of Poverty Measures

Measures of poverty necessarily refer to a segment of the population of a country or region, *i.e.*, that segment which has been measured as below the poverty line, or to be in poverty. Income or expenditure data for the households above the poverty line carry no weight in the calculations. This is the most important difference between measures of poverty on the one hand, and measures of either welfare or distribution on the other.

Ravallion draws a useful distinction between these types of indicators[10]:

- exclusive measures - attach zero weight to standards of living above a certain line, *e.g.*, poverty measure
- inclusive measures - give a positive weight to standards of living throughout the distribution, *e.g.*, the overall average per capita income

Poverty measures cannot be inferred from national income aggregates in the way that many welfare measures based on income or expenditure can be calculated. This is a characteristic that poverty measures share with measures of distribution. The estimation of poverty or distribution measures requires either household-level data (a typical primary source is a household survey) or tabulated grouped data.

Specific Measures of Poverty

There are three measures commonly used:[11]

- the headcount index
- the poverty gap index
- the squared poverty gap index

The headcount index is simply stated as the percentage of the population (or the percentage of households) which falls under the poverty line. Define **?** as the number of individuals or households below the poverty line and **n** as the total population (of either individuals or households). The headcount index is then expressed as:

$$HC = \frac{r}{n}$$

Of course, **?** itself is an indicator of poverty, but not in percentage terms.

A commonly cited flaw of the headcount index is its failure to estimate the depth or severity of poverty. For households higher than the poverty line, the welfare or income of households carries no weight in the calculation of the headcount index. However, the relative status of households below the poverty line carries no weight, either, in the headcount index. All poor households are given equal "weight" by this measure.

For example, if the worst-off households drop even further below the poverty line, the headcount index is unchanged. The problem of poverty, however, has inarguably increased in one of its dimensions. If the analyst is working with two or more poverty lines for headcount indexes (*e.g.*, the "lower" and "upper" poverty lines prepared by the World Bank), then this change will be observed only imprecisely.

Concerns about measuring the depth of poverty generated the design of the poverty gap index. The poverty gap index is defined as the headcount index multiplied by the gap between the mean income of the poor and the poverty line.[12] The mean income gap is expressed as the ratio of the mean income gap of the poor to the poverty line, *e.g.*, if mean income is three hundred pesos and the poverty line is four hundred pesos, then the gap would be 0.25.

Unlike the headcount index, welfare changes to a single household, or to a class of households among the poor, are reflected by changes in the poverty gap index.

There are several parts to the calculation of the poverty gap index. Because the headcount index is part of the formula for estimating the poverty gap index, that estimate must be obtained. Data must be available to array the income or expenditures in ascending order. We can define the lowest value of income or expenditure as $\boldsymbol{y}_1$ and the poverty line as ***l***. The poverty gap index (PG) is then measured as:

$$PG = \frac{1}{n}\sum_{i=1}^{r}\left[\frac{l - y_i}{l}\right]$$

Note that for incomes or expenditures above the poverty line, y_i is not included in the summation. The poverty gap ratio represents the mean proportionate poverty gap across the population. (The population has already been defined as ***n***.)

The right-hand term in the equation for PG includes the calculation of the "total income gap" for those in poverty. This is the figure summed up for all poor households. If this number is divided by the number of poor, ***?***, then the "mean income gap" is derived. Why not just use the "mean income gap" as the poverty indicator? As Ravallion points out, this is flawed.[13] If economic changes push the least poor above the poverty line -- and leave everyone below unchanged -- then the "mean income gap" will increase. The "gap" of the least poor was "closed," and dropped out of the calculation. The situation will appear to have worsened, even though the only change is that some households got lifted out of poverty.

By multiplying the "mean income gap" by either ***?*** or the headcount index, the poverty gap ratio defined as PG signals the improvement. In the definition of PG stated above, the "mean income gap" is multiplied by the headcount index (so the term ***?*** cancels out), which is equal to simply dividing the "total income gap" by the population, ***n***.

The "mean income gap" multiplied by ***?*** is a measure of the transfer that would bring every poor household up to the poverty line, *i.e.*, this is simply the "total income gap." This transfer would bring every poor household up to the poverty line assuming, of course that it could perfectly targeted. By contrast, if transfers cannot be targeted at all, then the transfer needed to bring every poor person up to the poverty line is simply the poverty line income or expenditure multiplied by the number of people in the entire population. As Ravallion notes, these can be interpreted as the minimum and maximum costs of eliminating poverty.[14]

In this context of eliminating poverty, the notion of "transfer" can be broadly interpreted. It could represent a payment from the public budget, or it can be an increase in income induced by economic growth. It can be expressed as an amount of money income, or as a percentage of GDP.

With the poverty gap index, the problem of not knowing the depth of poverty is eliminated. The measure PG is sensitive to changes in welfare or income experienced by poor households. However, the poverty gap index remains insensitive to changes in distribution ***within*** the universe of poor households. For example, if the income of the worst-off poor households goes up while the income of the least-poor falls by an equivalent amount, then the poverty gap index is unchanged. (The number of poor is also unchanged, so the headcount index remains the same.)

It seems that most analysts would agree that the change in the distribution described should be taken as a "lessening" of poverty by some measure. Consider the "depth" of poverty, as already used, as the typical distance from actual income to the poverty line for a poor household. The concept of depth refers to only one point, *i.e.*, the average distance from the poverty line. **A new concept -- the Aseverity@ of poverty -- refers to the distribution of the points around that single average point.** The severity of poverty can thought as taking into account the extreme, or maximum, distance from actual income to the poverty line for the worst-off households.

In other words, a measure designed to capture the "severity" of poverty would give more weight to the poverty of worst-off households than to those households just below the poverty line itself.

A common measure to capture this notion of "severity" is the squared poverty gap index (SPG).[15] The insight behind the SPG is that the worst-off poor are proportionately more "impoverished" than the group of poor near the poverty line. If you think of the "minimally acceptable" consumption level that the poverty line represents, a household near the poverty line will be forced to forego the "least necessary" of those goods and services. Households farther away from the poverty line are forced to forego the same goods and services, as well other goods and services that would be ranked -- dollar for dollar -- as having increasing value to their utility maximization.

The SPG is defined as:

$$SPG = \frac{1}{n}\sum_{i=1}^{r}\left[\frac{l - y_i}{l}\right]^2$$

with the variables defined as above in the poverty gap index. It gives more weight to the poverty gap of the worst-off households.

Consider two households, one of which is ten percent below the poverty line and the

other is fifty percent short. In the poverty gap measure (unsquared), the gap faced by the worse-off household carries fives time the weight of the other. In the squared poverty gap measure, the worse-off household has a weight of twenty-five times the other.

For example, assume that the poverty line is five hundred pesos, and every poor household is either at two hundred pesos or three hundred pesos. If all the two-hundred-peso households drop down to a hundred-peso level of income and the three-hundred-peso households rise up to a four hundred-peso level, what happens to the poverty measures? (The groups have equal numbers of households in this simplified example.)

< the number of poor is unchanged, so the headcount index is constant;
< the average poverty gap, or depth of poverty, is unchanged, so PG is constant;
< the "extreme" poverty gap, or the severity of poverty, does change, and the squared poverty gap registers a rise in poverty.

Although the SPG provides an estimate of the severity of poverty, changes in the index might be hard to interpret. Ravallion suggests that the SPG can be best thought of as having two components or sources of change:[16]

< the poverty gap, or depth of poverty itself; and
< 'inequality' among the poor.

In the case of perfect equality among the poor, the squared poverty gap index becomes identical to the poverty gap index.

The following table summarizes the three commonly-used measures of poverty:

Indicator	**Answers the question:**
Headcount Index	How many individuals or households are poor?
Poverty Gap Index	What is the depth of their poverty?
Squared Poverty Gap Index	What is the severity of their poverty?

How do these three indicators stack up in terms of their availability? Traditional number-crunching institutions might not always go this route, but Ravallion argues that primary data sources are not to blame. From his 1992 monograph (page 43):

> It is sometimes said that one rarely has access to the data needed to estimate the more "sophisticated" poverty measures....and so the best one can typically do is estimate the head-count index. This is not true. I have never seen a data set which would only permit estimation of the head-count index; indeed, I have never seen one for which the marginal cost of estimating the poverty gap....over the head-count index is anything but negligible with even modest computing resources.

The estimates can be generated whether the data set has individual or unit records (as is common

with a household survey data set), or tabulated grouped data.

A Generalized Form for Poverty Measures

Before moving on to commonly-used measures of welfare, a key point should be made about the three poverty measures heretofore defined: they represent a general class of measures and are representable in a single generalized form.

The Foster-Greer-Thorbecke Poverty Measures[17], as they are called, can be expressed as:

$$P_a = \frac{1}{n} \sum_{i=1}^{r} [\frac{l - y_i}{l}]^a$$

This definition closely resembles either the poverty gap or the squared poverty gap measured outlined above. The poverty measure, P_a, has a parameter a which determines which of the measures is obtained by the calculations. When a = 0, the headcount index is obtained. (Any quantity raised to the 0 power equals 1; the expression sums up ? data points.) When a = 1, the poverty gap index is obtained; when a = 2, the squared poverty gap index. To express more or less sensitivity to the severity of poverty, values of a greater than 1 and other than 2 can be inserted.

The UNDP's Non-income Measure of Poverty

The United Nations Development Program (UNDP) introduced a "Human Poverty Index"in 1996.[18] The HPI measures deprivation in basic human development by calculating the percentage of the population that fails to meet selected standards. These standards include: the percentage of people whose life expectancy is less than forty years, the adult illiteracy rate, the percentage of people without access to health care or safe drinking water, and the percentage of children under five who are measured as underweight. The HPI is the average of those percentages, with a higher score indicating a larger poverty problem.

Welfare

The set of "welfare" indicators is much broader than of poverty measures. As noted above, welfare measures do not set out a marker and then only assess the standard of living of the population on one side of the line. Welfare indicators attach a weight to all individuals or households in the distribution.

Most basically, there are three types of welfare indicators: income, non-income, and a composite involving both income and non-income. Consider the following list of welfare indicators, which is by no means exhaustive:

< Per capita income - usually expressed in U.S. dollars for comparability
< Infant mortality rate - expressed in number of deaths per thousand live births
< Human Development Index (HDI) - a composite of income and non-income indicators
< Physical Quality of Life Index (PQLI) - a composite of only non-income indicators

The inclusion of per capita income in this list should be taken as representative of a group of similar measures, including average household income or per capita consumption expenditure. The inclusion of the infant mortality rate is representative of an even longer list of non-income measures, including life expectancy in years, maternal mortality rates in deaths per 100,000 births, immunization rates, literacy rates, school enrolment ratios, percentage of population with access to safe drinking water, and per capita caloric consumption.

There are ways in which non-income welfare indicators can be expressed in a manner analogous to poverty indicators. Obviously, indicators such as the infant mortality rate or literacy rate, when calculated for only those households marked below the poverty line, are both poverty and welfare indicators. Alternatively, those non-income measures could be calculated on the basis of 'what percentage of the population falls below a certain benchmark or standard?'

The key point is that welfare indicators measure a standard of living along one or more dimensions for the society or region at large, while poverty indicators proceed from a segmentation of that group and measure the standard of living for that group exclusively.

The calculation of indicators such as per capita income (total income divided by the population) or infant mortality rates (total deaths before age 1 divided by the number of thousand live births) is straightforward. The rest of this section will present the methods used to calculate two commonly-used composite indicators, the HDI and the PQLI.[19]

***Human Development Index* (HDI)**

The United Nations Development Program (UNDP) publishes an annual edition of the HDI in its *Human Development Report* (Oxford University Press, various years). There are three categories of indicators in this composite measure:

- longevity - measured by life expectancy at birth
- educational attainment - a combination of literacy rates (b weight) and school enrolment ratios (a)
- standard of living - real GDP per capita in ppp-adjusted dollars[20]

The intuition of the HDI is to take minimum and maximum values for each of the four parts (educational attainment has two parts), and then calculate where each country lies along that spectrum. This is the general calculation for each part:

$$INDEX = \frac{Actual\ value\ of\ indicator - Minimum\ value\ of\ indicator}{Maximum\ value\ of\ indicator - Minimum\ value\ of\ indicator}$$

The markers for minimum and maximum values are given in the table on the following page.

	Minimum	**Maximum**
Life expectancy at birth	25 years	85 years
Adult literacy rate	0 percent	100 percent
School enrolment ratio	0 percent	100 percent
Real GDP per capita	PPP$ 100	PPP$ 40,000

For example, when life expectancy in a country is 55 years, the HDI formula yields a value of 0.50 for that component. This can be interpreted as saying that people in that country have attained only half the life expectancy that they should or could attain.

To calculate the “educational attainment” component, the formula is applied to both the literacy rate and the enrolment ratio, yielding two numbers that range from zero to one.

Then the formula

$$Educational\ attainment\ =\ \frac{2}{3}\,literacy\ index + \frac{1}{3}\,enrolment\ index$$

is applied, which also yields a number ranging from zero to one.

To calculate the "standard of living" component, the process is not as straightforward. An average of world income (measured as PPP$ 5,711 for 1995) is taken as a threshold income level, $\tilde{y}$, and all income above that line is discounted using Atkinson's utility of income formula. This formula assigns a declining marginal utility of income to all income levels above the threshold.

The upshot is that a "discounted value of PPP$ income" enters the HDI formula, not the estimate of income itself. For countries below the threshold income, no discounting is carried out.[21] After applying the HDI formula, another number between zero and one is generated.

The next step is to simply take the arithmetic average of the three index numbers. Each of the three components is given one-third weight. In turn, the HDI itself for each country lies between zero and one. Countries with the high marks near one are said to have attained most of the human development that can be expected (in terms of this measure), while countries with low marks closer to zero are found wanting.

***Physical Quality of Life Index* (PQLI)**

The PQLI has generated far less attention than the HDI, but represented an early attempt to focus development efforts on non-income dimensions.[22]

The PQLI is calculated very much in the spirit of the HDI, except that different components are used and no income measure is included. The three components in the PQLI (each with equal weight) are:

- life expectancy at age one
- infant mortality rate
- literacy rate

Unlike the HDI, which uses life expectancy at birth, the PQLI uses life expectancy at age one. This is done in order to avoid "double-counting" the impact of the infant mortality rate (not included in the HDI).

The three components are "normalized" to put them into common units in a manner analogous to that of the HDI. Because the literacy rate is already expressed as a number between zero and hundred, it is not 'normalized." The minimum-and-maximum-value approach used by the HDI is used for the life expectancy and infant mortality rates in the PQLI. For the minimum, the worst value attained by any country since 1950 is used. For the maximum, a "guesstimate"

about the best value that is possible to attain by the year 2000 is used. The PQLI formula is otherwise identical to that of the HDI formula, except that the values are multiplied by 100 in order to put them in percentage terms.

Once the three components are calculated and in "normalized" terms, the three numbers, each of which range from zero to one hundred, are added and divided by three to yield the PQLI. This number, in turn, ranges from zero (the worst possible score) to one hundred (the best possible score).

Uses of the HDI and the PQLI[23]

There are many examples of analysis which have used the HDI and its methods. Variations or modifications to the HDI have also been numerous, including several by the UNDP itself. This section is a quick summary of some examples.

Using the HDI to put countries in a matrix. Data indicate that income and social indicators are correlated, but that the relationship is imperfect. As noted above, nearly half of the countries in the HDI move up or down ten or more spaces in rank when the income-only measure is compared with the income-and-social-indicators measure. Fifteen percent of the countries move up or down twenty or more spaces.

As a means of putting focus on countries for further analysis, to say nothing of policy reforms, the HDI can be used to identify the exceptions to the Aincome-and-social-indicators@ tendency. Consider this simple matrix:

		Per Capita Income	
		High	Low
Social Indicators	Good	Group A (*e.g.*, OECD)	Group B (*e.g.*, Sri Lanka)
	Poor	Group C (*e.g.*, Mauritius)	Group D (*e.g.*, Bangladesh)

This is the two-by-two example, of course, and it could be expanded to suit one's tastes. Sri Lanka, as a low-income country that has achieved astounding progress on its social indicators, is the commonplace example. A country like Mauritius, with rapid economic growth in recent years, has lagged on its longevity and knowledge measures. That divergence, and others like it, merits some attention. The underlying point is longer lives, better health, and more education do not arise *ipso facto* from income growth by countries, but they rely upon a constellation of institutions, cultural changes, and policies themselves.

Constructing HDI's at the sub-national level as a starting point for policy dialogue. There are several examples of this approach, including "A Human Development Index for the Dalit

Child in India" by Bruce P. Corrie in the March 1995 issue of *Social Indicators Research*, and "A Human Development Index for the Black Child in the United States" by the same author in the January 1994 *Challenge*. The overriding purpose of the HDI is to bring attention to social and economic problems. Application of the method to sub-national data is a natural extension of the method.

Tracking broad trends over longer time periods. Intertemporal comparisons were a bit dicey with earlier versions of the HDI. Because each year's worth of data was used to set the minimum and maximum points, the "goalposts" were moving each year. However, the UNDP has adjusted the HDI so that the minimum and maximum markers are fixed. (Note: there is still some residual variability in the treatment of diminishing returns to income.) Now comparisons over time can be made with integrity. Most data are available for the past thirty years or so, and the HDI can be "back-calculated" even if the UNDP does not publish such estimates.

It is noted that the HDI does not vary much year-to-year, but this is a feature of life rather than the concept of the index. The more broadly inclusive the economic or social aggregate, the less likely that it will vary greatly on an annual basis.

The Gender-related Development Index (GDI). This comes out of the UNDP itself, making its debut in 1995. The GDI measures the same capabilities as the HDI, but it is adjusted for different levels of achievement for women and men. In order to calculate the GDI, a method of penalizing inequality is used. A country's GDI will be lower -- relative to its HDI score -- where there is greater gender inequality.

The Gender Empowerment Measure (GEM). Also from the UNDP starting in 1995, the GEM is designed to measure the extent to which women and men are equal in their empowerment to participate in economic and social life, and to affect decision-making. The first step is to calculate indexes for the gender balances in managerial, administrative, professional, and technical occupations, as well as in parliamentary representation. This is adjusted by the method of "population-weighted averaging." Second, the gender balance in the distribution of earned income is calculated. This is adjusted by the gender balance in the "economically active population," rather than the population as a whole. (For further details, see the technical notes to the 1996 *Human Development Report*.)

The Capability Poverty Measure (CPM). Also from the UNDP itself, the CPM is a simple index which calculates the percentage of the national population which does not have adequate capability in three dimensions of human development: a healthy and well-nourished life, the ability for safe and healthy reproduction, and being knowledgeable and literate. Three indicators are used: percentage of children under five years old who are underweight, percentage of births unattended by trained health personnel, and the percentage of women aged fifteen years or older who are illiterate.

In the case of the HDI, attention is paid to the average level of capabilities in a country; in the case of the CPM, attention is given to the share of people who lack those capabilities. (See the technical notes in the *Human Development Report* for further details.)

Conclusion on Welfare Measures. Leaving any debates over the HDI itself aside, it is worth noting that intense focus on the HDI does an injustice to the *Human Development Report* and its plenitude of data tables. The usefulness of the *HDR* as a development document is best seen in the wealth of indicators that it assembles in one place, rather than in the headline-grabbing country ranking.

Distribution

Measures of distribution are nearly always based on the dimensions of income, consumption expenditure, or assets. Unlike the group of welfare indicators considered in this technical report, measures of distribution utilize data at the individual or household level, rather than simply take country or regional averages. Unlike the group of poverty indicators considered in this chapter, measures of distribution do not necessarily exclude data with respect to the non-poor.

Common measures of distribution include:

< the Gini coefficient
< the ratio of income for the highest twenty percent to the lowest forty percent of individuals or households

Calculation of these measures relies upon access to household-level data or tabulated group data.

The Gini Coefficient

This is the most commonly-cited measure of the inequality of the distribution of income in a society. The value of the Gini coefficient will vary from zero (perfect equality) to one (perfect inequality). The simplest way to present the calculation of the Gini is with a graph.

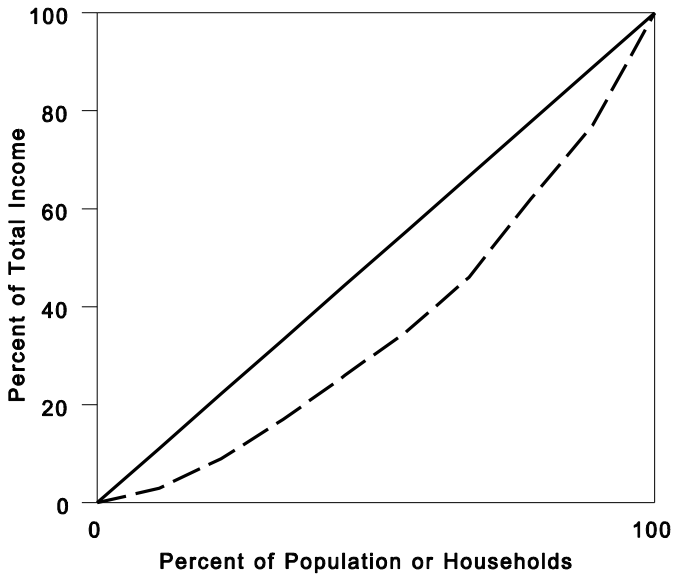

First, data on income for each household are arrayed so that the lowest is first and the others are in ascending order. To generate the dashed line in the Gini graph (above), the question is asked: what percentage of total income is earned by the x^{th} percentage of the population?

For example, to represent the lowest ten percent of households, add up their income and calculate it as a percentage of total income. Because of the way in which the data are arrayed, this percentage of total income will be no greater than ten percent.

Each point on the dashed line represents a household, placed by virtue of its rank in income level. Each point on the dashed line represents the total or cumulative percent share of income for every household up to and including that particular household.

By definition, the dashed line must be at zero percent of total income when it is at zero percent of total population; and at one hundred percent of total income when at one hundred percent of total population. This is because, no matter what the distribution of income, one hundred percent of the population has one hundred percent of the income.

Consider the shape of the dashed line. If incomes are very much unequally distributed, the dashed line will stretch out to the right. (It will not shift, in the sense that it must intersect zero and one hundred.) For example, the less income earned by the bottom half of households, the more that is earned by the top half – this is always true for any given level of total income. In that case, the dashed line is lower up to the point of fifty percent of the population, but then rises more sharply for the rest of the line.

The solid line in the Gini graph represents perfect equality of incomes. For example, if all incomes are equal, and households are ranked by income level, then fifty percent of the population will have fifty percent of the income. All down the line, the percentages of total income and population will be equal.

The dashed and solid lines form a shape that looks like an orange wedge. The Gini coefficient is defined as the ratio of the area in that orange wedge to the total area of the triangle formed by the solid line. This coefficient will be a number between zero (*i.e.*, the dashed line equals the solid line) to one (*i.e.*, the dashed line equals the frame of the graph).

There are at least two ways to calculate the Gini coefficient; see the End Notes to Part I for further details.[24]

Relative Income Shares

This indicator is calculated in a straightforward manner. Using tabulated grouped data -- often available from many sources -- the ratio of income shares between different groups is compared. It is common to use the income share represented by the top twenty percent of households, or the top quintile. The ratio of that share is taken with respect to the income share represented by the bottom forty percent.

The minimum value for this ratio (T20/B40) is 0.5, which represents the case of incomes being perfectly equal across all households. (In that case, the income share of each quintile is twenty.) In practice, however, the ratio T20/40B appears to be at least greater than one, and typically between two and five.

For example, using data in the 1996 *Social Indicators of Development* from the World Bank, the quintile income shares are provided for eighteen countries for the year 1992. The ratio of top to bottom income shares ranges from 1.13 in the Slovak Republic to 6.11 in Kenya. The median value for the eighteen countries is 2.17. (All of these eighteen countries are either low- or middle-income countries.)

In the Bank's annual *World Development Report*, income share data are available for the majority of countries -- but only for years a decade or more ago. Data for the United States, for example, is from 1985 and the T20/B40 ratio is 2.67. This compares with 2.74 for Australia (also 1985) and 1.74 for Sweden (1981).[25]

Incorporating A Distributional Measure into a Welfare Indicator

One way in which to make a welfare indicator, *e.g.*, per capita income, sensitive to distributional considerations is the social welfare measure developed by Atkinson in a constant-elasticity-of-substitution form.[26]

$$SW_e = \frac{1}{n}\sum_{i=1}^{n} y_i^{1-e}$$

There are n individuals or households, each with income y_i. The measure of social welfare has a nonnegative parameter **e** which is the elasticity of marginal social welfare. This is also understood as expressing the extent of inequality aversion. Increasing values of **e** are associated with greater aversion to inequality of income.

End Notes

1. Ravallion states that "'poverty' can be said to exist in a given society when one or more persons do not attain a level of material well-being deemed to constitute a reasonable minimum by the standards of that society." (*Poverty Comparisons: A Guide to Concepts and Methods*, LSMS Working Paper Number 88, 1992, page 4.) According to *The MIT Dictionary of Modern Economics*, fourth edition (David W. Pearce, editor, 1992), poverty can be either relative to the standards of society, or absolute and reflective of a minimum condition for survival. In practice, most data is really both- it is an absolute line drawn at the point that reflects a relative standard for a given society. For example, in a recent World Bank study, *Five Small Open Economies*, edited by Ronald Findlay and Stanislaw Wellisz (1993), the case studies are tied together with comparisons, including a "fixed" poverty line for each country. The "fixed" lines, however, are different for each country.

In this report, only the standard income- or expenditure-based measures of poverty are presented. Other approaches are both possible and often highly valuable. For example, an anthropological method can be applied which would examine the cultural characteristics of poverty. Two studies which take up this approach are: John Iliffe, *The African Poor: A History*, Cambridge University Press, 1987; and Willem Van Schendel, *Three Deltas: Accumulation and Poverty in Rural Burma, Bengal, and South India*, Sage Publications, 1991.

2. The choice of variable upon which to base a poverty line is generally between income and expenditure. Gaurav Datt and Peter Lanjouvw of the World Bank outlined some of the arguments for and against each choice. ("Measurement of Poverty and Its Impact," paper prepared for the USAID Training Workshop, September 1995.) Income, when measured at the household level, poses several challenges. Many households, particularly in rural areas, have multiple sources of income and a total income measure might not be readily available. The seasonality and long-term cycles of agricultural income can pose problems for accuracy of measurement. Self-employment and the informal sector are difficult to capture with traditional data collection methods. For these and other reasons, there are advantages to using a consumption measure as the poverty line. For example, consumption tends to fluctuate less than income and allows for a better measure of a household's typical status. It eliminates the need to aggregate diverse sources of income. However, there are also difficulties with the choice of consumption measures. Controlling for qualitative differences in goods and services is hard to do; the choice of the consumption "basket" is subjective; and the data are usually in money rather than volume terms. With respect to transition economies, the appropriate goods and services in the appropriate consumption basket are often undergoing rapid change.

3. Martin Ravallion and Binayak Sen, "When Method Matters: Monitoring Poverty in Bangladesh," *Economic Development and Cultural Change*, July 1996 (44:4), pages 761-792.

4. Not requiring price data is a strong reason to favor an FEI approach- and underscores the biggest obstacle to generating poverty estimates: the dearth of reliable data. Lacking data, and

perhaps a consensus on the "minimum" basket of consumption goods, a possible approach is to set a relative poverty line for a society "as is," *e.g.*, thirty percent of median income. However, as Grant Scobie argues in his *Macroeconomic Adjustment and the Poor: Toward a Research Policy* (Cornell Food and Nutrition Policy Program, Monograph Number 1, April 1989, pages 70-74), "relative measures" of poverty are not really about poverty at all, but about the separate problem of inequality. According to Scobie, there is an inherent "absoluteness" in the concept of poverty. This point is well taken, but the seasoned analyst does well to quibble over whether the problems of poverty and inequality are indeed completely "separate."

5. In the United States, the seminal work on how to draw a poverty line includes Mollie Orshansky's "Counting the Poor: Another Look at the Poverty Profile," *Social Security Bulletin*, volume 28, 1965, pages 3-29. Orshansky used United States Department of Agriculture data on the cost of a minimally adequate food budget, made adjustments for family size and composition, and then multiplied that number by three. Her reasoning was that food typically consumed one-third of household budgets at that income level, a factor which found support in some surveys. Even though Orshansky's work represented a major improvement over previous efforts, it remains open to several criticisms. Nonetheless, the basic Orshansky approach remains behind official U.S. poverty statistics. A recent and comprehensive look at the measurement issues in the United States is Patricia Ruggles's *Drawing the Line: Alternative Poverty Measures and Their Implications for Public Policy*, Urban Institute Press, 1990. Ruggles argues that a complete updating of what "consumption needs" are deemed minimally adequate is needed every decade or so (page 2).

6. When working with estimates of income or expenditure in different countries, the first problem for comparative analysis is that the data are in different units or local currencies. One approach is to convert all data into U.S. dollars, using the prevailing exchange rate on the currency markets. However, a substantial body of studies has shown that the foreign exchange markets do not typically produce exchange rates that "equilibrate" prices in different countries. If they did, economists refer to that state of affairs as the "purchasing power parity" (ppp) exchange rate. For example, if a loaf of bread cost one pound in London and five pesetas in Madrid, the purchasing power parity exchange rate (assuming that folks do live by bread alone) would be £1:P5. That would imply that if you took your pounds, went to Spain, and exchanged them for pesetas, you would not notice any differences in prices.

In most countries and at most times, you do in fact notice price differences, or differences in purchasing power. The United Nations launched an International Comparisons Project and the University of Pennsylvania developed the "Penn World Tables" in order to calculate an exchange rate which does equilibrate prices in different countries. When the ppp-adjusted exchange rates are used to convert estimates originally expressed in local currencies, the resultant units are typically known as "international dollars," or simply dollars.

7. The World Bank, *Implementing the World Bank's Strategy to Reduce Poverty*, 1993, p. 5.

8. How important is the distinction between "poverty" as measured by this upper poverty line, and "extreme poverty" as measured by the lower poverty line? In the 1990 *World Development*

Report, the World Bank provided an example of a poverty study for Indonesia in 1981. At issue was the impact of higher food staple prices on the poor. Assuming that the higher price of rice was passed on to the producers, one study determined that the poor would benefit. On average, the poor were net producers of food staples.

In other words, boosting the price of food would reduce the number of poor so counted by using the upper poverty line. However, the study also determined that the "poorest of the poor," *i.e.*, those in extreme poverty, were primarily landless and not net producers of food staples. Therefore, higher food prices by themselves would hurt their interests and raise the number of people counted as poor using the lower poverty line. (page 28) The divergence in these two measures is probably not unique to Indonesia in the early 1980s.

What is important here is not so much finding the right place to draw those poverty lines, but achieving a deeper understanding of the socioeconomic groups within an society in order to understand the impact of policies and trends. With poverty analysis, it is important to read between the lines.

9. Shaohua Chen, Gaurav Datt, and Martin Ravallion, "Is Poverty Increasing in the Developing World?" World Bank *Working Papers*, WPS 1146, June 1993, table 4, page 28.

10. Martin Ravallion, "Measuring Social Welfare With and Without Poverty Lines," *American Economic Review*, (84) 2, May 1994, pages 359-364.

11. Much of the work on poverty concepts and measurement has been done by economists at the World Bank, including Martin Ravallion, Gaurav Datt, Shubham Chaudhuri, Peter Lanjouw, Benu Bidani, Dominique van de Walle, and Shaohua Chen. See the References section for citations. Ravallion's *Poverty Comparisons* is probably the single most comprehensive reference. It appeared as a World Bank LSMS Working Paper (Number 88) in 1992, as well as in book form in 1994 (New York: Harwood).

12. This section borrows heavily from Ravallion's *Poverty Comparisons* (LSMS Working Paper, *ibid.*), particularly pages 36-40.

13. See page 37, *ibid*.

14. We defined the poverty gap measure as:

$$PG = \frac{1}{n}\sum_{i=1}^{r}\left[\frac{l - y_i}{l}\right]$$

where ***?*** is the number of poor people, ***l*** is the poverty line in either income or expenditure terms, y_i is individual incomes arrayed in ascending order, and ***n*** is the total population. If we calculate the mean income of the poor, $y^{?}$, then the "income gap ratio" (*I)* is defined as:

$$I = \frac{l - y^r}{l}$$

Multiply that "income gap ratio," which is the mean distance below the poverty line as a percent of the poverty line, times the number of poor, ***?***, and you have the "minimum" cost of eliminating poverty. This is the minimum because it assumes that the policies and/or transfers can be perfectly targeted. If all benefits reach the poor and only the poor, the poverty gap will be closed by a transfer of ***?*** multiplied by *I*.

If transfers cannot be targeted at all, then the cost of eliminating poverty is simply the product of ***l*** and ***n***, because the policy maker must ensure that everyone is at least as well-off as the poverty line. If we consider the ratio of the minimum to the maximum cost of eliminating poverty:

$$\frac{r \times I}{l \times n}$$

We can substitute the definition of *I* and rearrange terms to yield the equivalent:

$$\frac{r}{n} \times l \left[\frac{l - y^r}{l} \right]$$

By simplification and recognition that ***?*** times the term (***l*** - ***y***^{I}) is the same as:

$$\sum_{i=1}^{r} \frac{l - y_i}{l}$$

we can write the ratio of minimum to maximum costs as:

$$\frac{1}{n} \sum_{i=1}^{r} \frac{l - y_i}{l}$$

which is the poverty gap (PG) as defined in the text.

Ravallion notes that the poverty gap is therefore an indicator of the implicit bonus from successfully targeting anti-poverty transfers and policies.

15. This concept has a long background in the literature. Among the key articles are: Amartya Sen, "Poverty: An Ordinal Approach to Measurement," *Econometrica*, (44) 1976, pages 219-231; Nanak Kakwani, "On A Class of Poverty Measures," *Econometrica* (48) 1980, pages 437-446; Amartya Sen, *Poverty and Famines: An Essay on Entitlement and Deprivation*, OUP, 1981; James Foster, J. Greer, and E. Thorbecke, "A Class of Decomposable Poverty Measures," *Econometrica* (52) 1984, pages 761-765; Ravallion, *Poverty Comparisons*, *op.cit.*; and Soniya

Carvalho and Howard White, "Indicators for Monitoring Poverty Reduction," *World Bank Discussion Papers,* Number 254, 1995.

16. Ravallion, *Poverty Comparisons*, *op. cit.*, page 39.

17. Foster, *et. al.*, "A Class of Decomposable Poverty Measures," *op. cit.* The measures are applied in Martin Ravallion, "Measuring Social Welfare With and Without Poverty Lines," *op. cit.*

18. See the UNDP's *Human Development Report* for either 1996 or 1997 for further details.

19. The HDI and PQLI are not alone among published composite indicators of welfare. Population Action International in Washington, DC published *The International Human Suffering Index* in 1992. The Cato Institute, *et. al.*, published an *Economic Freedom Index*, authored by James Gwartney, *et. al.*, in 1996 with a revised edition expected in 1997.

20. See Technical Note 1 in the 1996 *Human Development Report* (UNDP).

21. Atkinson's utility of income of formula can be summarized as:

$W(y) = y_i$ when $y_i <$ threshold income y^*

$W(y) = y^* + 2[(y_i - y^*)^{1/2}]$ when $y^* = y_i = 2y^*$

$W(y) = y^* + 2(y^{*1/2}) + 3[(y_i - 2y^*)^{1/3}]$ when $2y^* = y_i = 3y^*$

As country *i* per capita income rises to greater multiples of the threshold income, the formula is expanded in the same manner. This method is used to determine that the maximum real income of PPP$ 40,000 is equivalent to a discounted income of PPP$ 6,040.

Examples are given in the technical notes to the 1996 *Human Development Report*. Earlier versions of the report should not be used for instructions on methods because the HDI has been fine-tuned just about every year it has been published.

22. Morris David Morris, *Measuring the Condition of the World's Poor: The Physical Quality of Life Index*, Pergamen Press, 1979. Morris intended to measure countries on their performance in satisfying selected and specific life-serving social characteristics. Unlike the other indicators in this report, the PQLI is not directly available from any official source. However, the estimates can be easily calculated from existing data.

23. For further details on the uses and variations of the HDI, a short reference list includes: Desai, Meghnad, "Human Development: Concepts and Measurement," *European Economic Review*, April 1991, pages 350-357; Lüchers, Guido and Lukas Menkhoff, 'Human Development as Statistical Artifact," *World Development*, August 1996, page 1385-1392; Srinivasan, T. N., "Human Development: A New Paradigm or Reinvention of the Wheel?" *AEA Papers and*

Proceedings, May 1994, pages 238-243; Streeten, Paul, "Human Development: Means and Ends," *AEA Papers and Proceedings*, May 1994, pages 232-237; and, United Nations Development Program, *Human Development Report*, annual since 1990.

24. Gini coefficient formulas are given in R. M. Sundrum, *Income Distribution in Less Developed Countries*, Routledge, 1990, pages 46-70. Following Sundrum, one way to think of the Gini coefficient is as the relative mean difference in income, divided by twice the mean income μ. The equation:

$$G = \frac{1}{2m} \sum \left| x_i - x_j \right|$$

To calculate the relative mean difference, all actual differences between any two pairs of n households are summed up and then this sum is divided by n^2.

Alternatively, one can think of the dashed lines in Gini graphs as representing smooth functions of the form F(x). The Gini coefficient can be expressed as:

$$G = \frac{2}{m} \int_{8}^{\infty} x \left[F(x) - 1/2 \right] f(x) dx$$

You do the math.

See the book by Sundrum for a description of how to calculate the Gini coefficient when using tabulated grouped data.

25. The World Bank, *World Development Report 1996*, Oxford University Press, table 5 in the "Selected World Development Indicators" section.

26. This measure is described in Ravallion, "Measuring Social Welfare With and Without Poverty Lines," *op. cit.*, page 362. The original work was by Anthony Atkinson, "On the Measurement of Inequality," *Journal of Economic Theory*, September 1970, 2(3), pages 244-263.